

Requested Patent: WO9110998A1

Title: DATA STORAGE ;

Abstracted Patent: WO9110998 ;

Publication Date: 1991-07-25 ;

Inventor(s):

SIMMES MARK (GB); BRAMHALL PETER (GB); VAN MAREN DAVID (US) ;

Applicant(s): HEWLETT PACKARD LTD (GB) ;

Application Number: WO1991GB00082 19910118 ;

Priority Number(s): GB19900001334 19900119 ;

IPC Classification: G11B20/00 ; G11B20/12 ;

Equivalents: EP0464181 (WO9110998), JP5500878T

ABSTRACT:


A data storage method for writing compressed data organised in the form of records (CRn) to tape (10) characterised by inserting into the datastream ancillary information which is extra to the data compression process.



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁵ : G11B 20/00, 20/12</p>	<p>A1</p>	<p>(11) International Publication Number: WO 91/10998</p> <p>(43) International Publication Date: 25 July 1991 (25.07.91)</p>
<p>(21) International Application Number: PCT/GB91/00082</p> <p>(22) International Filing Date: 18 January 1991 (18.01.91)</p> <p>(30) Priority data: 9001334.3 19 January 1990 (19.01.90) GB</p> <p>(71) Applicant (for all designated States except US): HEWLETT-PACKARD LIMITED [GB/GB]; Cain Road, Bracknell, Berkshire RG12 1HN (GB).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only) : VAN MAREN, David [US/US]; Hewlett-Packard Company, Greeley Storage Division R&D, 700 71st Street, Greeley, CO 80634 (US). SIMMES, Mark [GB/GB]; BRAMHALL, Peter [GB/GB]; Hewlett-Packard Limited, Filton Road, Stoke Gifford, Bristol BS12 6OZ (GB).</p>		<p>(74) Agent: SMITH, Denise, Mary; Hewlett-Packard Limited., Intellectual Property Section, Building 2, Filton Road, Stoke Gifford, Bristol BS12 6OZ (GB).</p> <p>(81) Designated States: AT (European patent), BE (European patent), CH (European patent), DE (European patent), DK (European patent), ES (European patent), FR (European patent), GB (European patent), GR (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent), US.</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: DATA STORAGE



(57) Abstract

A data storage method for writing compressed data organised in the form of records (CR_n) to tape (10) characterised by inserting into the datastream ancillary information which is extra to the data compression process.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TC	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

DATA STORAGE

The present invention relates to the compression of user data and its storage on tape.

5 It is known to provide a tape drive having data compression capability (a DC drive) so that, as data arrives from a host, it is compressed before being written to tape thus increasing the tape storage capacity. DC drives are also able to read compressed data from tape and to
10 decompress the data before sending it to a host. It is also possible for a host to perform software compression and/or decompression of user data.

There is more than one type of data compression. For example, removing separation marks eg. designating records,
15 files etc. from the datastream and storing information regarding the positions of these marks in an index effectively compresses the user data. Another, quite different approach, is to compress user data by removing
20 redundancy in the data eg. by replacing user data words with codewords from which the original data can be recovered. It is the latter type which is being referred to in this specification when the words "data compression" or abbreviation DC is used.

According to the present invention we provide a data
25 storage method for writing compressed data organised in the form of records to tape characterised by inserting into the datastream ancillary information which is extra to the data compression process.

The ancillary information may comprise error checking
30 information. Furthermore, the ancillary information may comprise data separation information ie. information which could be used to separate the data later.

The aim of inserting this extra information into the datastream as part of a data compression algorithm is to
35 render the datastream particularly suitable for fast operation and easy checking of data error conditions. For example, codewords representing the uncompressed byte count

and/or a redundancy check could be inserted after an "end of record" codeword. These codewords could be utilised during error checking operations but could be skipped if they are not required or are inappropriate for particular tape drive.

5 The method preferably comprises writing the ancillary information to tape in uncompressed form. This is preferred so that the ancillary information is available to a non-DC tape drive.

10 The method may comprise inserting into the datastream ancillary information in association with one or more records.

 The method may comprise inserting into the datastream a header portion containing ancillary information relating to one or more records following the header portion.

15 The method may further comprise inserting into the datastream a trailer portion containing ancillary information relating to one or more records preceding the trailer portion.

20 Alternatively, or as well, the method may comprise organising data records into groups independently of the record structure of the data, and writing information regarding the records in a group to an index associated with the group.

25 In an embodiment to be described, the method comprises writing information to the group indices in terms of entities, where an entity comprises one or more records. In that embodiment, the method comprises writing ancillary information to a header associated with each entity.

30 The present invention also provides a storage device for compressing user data and writing compressed data to tape which is operable in accordance with a method as defined above.

35 Particular embodiments of the present invention will now be described, by way of example, with reference to the accompanying diagrammatic drawings in which:

 Figures A and B are diagrams relating to a data compression algorithm;

Figure 1 is a multi-part diagram illustrating a scheme for storing computer data where:

(a) is a diagram representing a sequence of data records and logical separation marks sent by a user (host) to data storage apparatus;

(b) and (c) are diagrams illustrating two different arrangements for storing the sequence of Figure 1 (a) on tape;

Figure 2 is a diagram of a group index;

Figures 3 and 3A are diagrams of general block access tables;

Figures 4 and 4A are diagrams of specific block access tables;

Figures 5 - 7 are diagrams of further schemes for storing computer data;

Figure 8 is a diagram illustrating possible valid entries for the block access table of a group.

Figures 9 and 10 are further diagrams of schemes for storing computer data;

Figure 11 is a diagram illustrating the main physical components of a tape deck which employs helical scanning and which forms part of the data storage apparatus embodying the invention;

Figure 12 is a diagrammatic representation of two data tracks recorded on tape using helical scanning;

Figure 13 is a diagrammatic representation of the format of a main data area of a data track recorded in accordance with the present data storage method;

Figure 14 is a diagrammatic representation of the format of a sub data area of a data track recorded in accordance with the present data storage method;

Figure 15 is a diagram showing for the present method, both the arrangement of data frames in groups within a data area of a tape and details of an index recorded within each group of frames;

Figure 16 is a block diagram of the main components of the data storage apparatus embodying the invention;

Figures 17 and 18 are block diagrams relating to the data compression processor;

Figure 19 is more detailed functional block diagram of a group processor of the data storage apparatus;

5 Figures 20A and 20B are flow charts of algorithms implemented by the drive apparatus in searching for a particular record on a tape.

Further information regarding data compression, including details of a specific DC algorithm will first be given.

The aim of a data compression process is to remove redundancy from data. One measure of compression efficiency is called "compression ratio" and is defined as:

15

$$\frac{\text{Length of uncompressed input}}{\text{Length of compressed output}}$$

This is a measure of the success of a data compression process. The larger the compression ratio, the greater the compression efficiency.

One way of performing data compression is by recognising and encoding patterns of input characters, ie. a substitutional method.

25 According to the LZW algorithm, as unique strings of input characters are found, they are entered into a dictionary and assigned numeric values. The dictionary is formed dynamically as the data is being compressed and is reconstructed from the data during decompression. Once a dictionary entry exists, subsequent occurrences of that entry within the datastream can be replaced by the numeric value or codeword. It should be noted that this algorithm is not limited to compressing ASCII text data. Its principles apply equally well to binary files, data bases, 30 imaging data, and so on.

Each dictionary entry consists of two items: (1) a unique string of data bytes that the algorithm has found

within the data, and (2) a codeword that represents this combination of bytes. The dictionary can contain up to 4096 entries. The first eight entries are reserved codewords that are used to flag and control specific conditions. The
5 next 256 entries contain the byte values 0 through 255. Some of these 256 entries are therefore codewords for the ASCII text characters. The remaining locations are linked-list entries that point to other dictionary locations and eventually terminate by pointing at one of the byte values
10 0 through 255. Using this linked-list data structure, the possible byte combinations can be anywhere from 2 bytes to 128 bytes long without requiring an excessively wide memory array to store them.

In a hardware implementation of the scheme which will
15 be more fully described later, the dictionary is built and stored in a bank of random-access memory (RAM) that is 23 bits wide. Each memory address can contain a byte value in the lower 8 bits, a codeword or pointer representing an entry in the next 12 bits, and three condition flags in the
20 upper 3 bits. The number of bits in the output byte stream used to represent a codeword ranges from 9 bits to 12 bits and corresponds to dictionary entries that range from 0 to 4095. During the dictionary building phase, until 512 entries are made into the dictionary 9-bits are used for
25 each codeword, after the 512th entry 10-bits are needed for the codewords, after the 1024th entry 11-bits are needed for the codewords, and for the final 2048 entries 12-bits are needed for the codewords. Once the dictionary is full, no further entries are built, and all subsequent codewords are
30 12 bits in length. The memory address for a given dictionary entry is determined by a complex operation performed on the entry value. Since the dictionary can contain 4096 entries, it would appear that 4K bytes of RAM is all that is needed to support a full dictionary. This is
35 in fact the case during decompression. However, during compression, more than 4K bytes of RAM is needed because of dictionary "collisions" that occur during the dictionary

building phase. This is when two different string character combinations map to the same location in the dictionary RAM and is a consequence of the finite resources in dictionary RAM and the complex process of dictionary building during
5 compression. When a dictionary collision occurs, the two colliding values are recalculated to two new locations and the original location is flagged as a collision site.

An important property of the algorithm is the coupling
10 between compression and decompression. These two operations are tied together both in the compression and decompression processes and in the packing and unpacking of codewords into a byte stream. The nature of the compression algorithm requires that the compression process and the decompression
15 process be synchronized. Stated differently, decompression cannot begin at an arbitrary point in the compressed data. It begins at the point where the dictionary is known to be empty or reset. This coupling provides one of the fundamental advantages of the algorithm, namely that the
20 dictionary is embedded in the codewords and does not need to be transferred with the compressed data. Similarly, the packing and unpacking process must be synchronized. Note that compressed data must be presented to the decompression hardware in the proper order.

25 Fig A is a simplified graphical depiction of the compression algorithm referred to above. This example shows an input data stream composed of the following characters: R I N T I N T I N. To follow the flow of the compression process, Fig A should be viewed from the top to the bottom,
30 starting at the left and proceeding to the right. It is assumed that the dictionary has been reset and initialized to contain the eight reserved codewords and the first 256 entries of 0 to 255 including codewords for all the ASCII characters.

35 The compression algorithm executes the following process with each byte in the data stream:

1. Get the input byte.

2. Search the dictionary with the current input sequence and, if there is a match, get another input byte and add it to the current sequence, remembering the largest sequence that matched.
 - 5 3. Repeat step 2 until no match is found.
 4. Build a new dictionary entry of the current "no match" sequence.
 5. Output the codeword for the largest sequence that matched.
- 10 In this example, the compression algorithm begins after the first R has been accepted by the compression engine. The input character R matches the character R that was placed in the dictionary during its initialization. Since there was a match, the DC engine accepts another byte,
- 15 this one being the character I. The sequence RI is now searched for in the dictionary but no match is found. Consequently, a new dictionary entry RI is built and the codeword for the largest matching sequence (i.e., the codeword for the character R) is output. The engine now
- 20 searches for I in the dictionary and finds a match just as it did with R. Another character is input (N) and a search begins for the sequence IN. Since IN does not match any entries, a new one is built and the codeword for the largest matching sequence (i.e., the codeword for the character I)
- 25 is output. This process continues with a search for the letter N. After N is found, the next character is input and the dictionary is searched for NT. Since this is not found, a dictionary entry for NT is built and the codeword for N is output. the same sequence occurs for the characters T and
- 30 I. A codeword for T is output and a dictionary entry is built for TI.

Up to this point, no compression has occurred, since there have been no multiple character matches. In actuality, the output stream has expanded slightly, since

35 four 8-bit characters have been replaced by four 9-bit codewords. (That represents a 32-bit to 36-bit expansion,

or a 1.125:1 compression ratio.) However, after the next character has been input, compression of the data begins.

At this point, the engine is searching for the IN sequence.

Since it finds a match, it accepts another character and

5 begins searching for INT. When it does not find a match, it

builds a dictionary entry for INT and outputs the previously

generated codeword for the sequence IN. Two 8-bit

characters have now been replaced by one 9-bit codeword for

a compression ratio of 16/9 or 1.778:1.

10 This process continues and again two characters are

replaced with a single codeword. The engine begins with a

T from the previous sequence and then accepts the next

character which is an I. It searches for the TI sequence

and finds a match, so another byte is input. Now the chip

15 is searching for the TIN sequence. No match is found, so a

TIN entry is built and the codeword for TI is output. This

sequence also exhibits the 1.778:1 compression ratio that

the IN sequence exhibited. The net compression ratio for

this string of 9 bytes is 1.143:1. This is not a

20 particularly large compression ratio because the example

consists of a very small number of bytes. With a larger

sample of data, more sequences of data are stored and larger

sequences of bytes are replaced by a single codeword. It is

possible to achieve compression ratios that range from 1:1

25 up to 110:1.

A simplified diagram of the decompression process is

shown in Fig B. This example uses the output of the

previous compression example as input. The decompression

process looks very similar to the compression process, but

30 the algorithm for decompression is less complicated than

that for compression, since it does not have to search for

the presence of a given dictionary entry. The coupling of

the two processes guarantees the existence of the

appropriate dictionary entries during decompression. The

35 algorithm simply uses the input codewords to look up the

byte sequence in the dictionary and then builds new entries

using the same rules that the compression algorithm uses. This is the only way that the decompression algorithm can recover the compressed data without a special dictionary sent with each data packet.

5 As in the compression example, it is assumed that the dictionary has been reset and initialized to contain the first 256 entries of 0 to 255. The decompression engine begins by accepting the codeword for R. It uses this codeword to look up the byte value R. This value is placed
10 on the last-in, first-out (LIFO) stack, waiting to be output from the chip. Since the R is one of the root codewords (one of the first 256 entries), the end of the list has been reached for this codeword. The output stack is then dumped from the chip. The engine then inputs the codeword for I
15 and uses it to look up the byte value I. Again, this value is a root codeword, so the output sequence for this codeword is completed and the byte value for I is popped from the output stack. At this point, a new dictionary entry is built using the last byte value that was pushed onto the
20 output stack (I) and the previous codeword (the codeword for R). Each entry is built in this manner and contains a byte value and a pointer to the next byte in the sequence (the previous codeword). A linked list is generated in this manner for each dictionary entry.

25 The next codeword is input (the codeword for N) and the process is repeated. This time an N is output and a new dictionary entry is built containing the byte value N and the codeword for I. The codeword for T is input, causing a T to be output and another dictionary entry to be built.
30 The next codeword that is input represents the byte sequence IN. The decompression engine uses this codeword to reference the second dictionary entry, which was generated earlier in this example. This entry contains the byte value N, which is placed on the output stack, and the pointer to
35 the codeword for I, which becomes the current codeword. This new codeword is used to find the next byte (I), which

is placed on the output stack. Since this is a root codeword, the look up process is complete and the output stack is dumped in reverse order, that is, I is output first, followed by N. The same process is repeated with the
5 next two codewords, resulting in the recovery of the original byte sequence R I N T I N T I N.

Two of the reserved codewords mentioned above which are inserted into the data stream during data compression are the RESET and FLUSH codewords. The RESET codeword
10 signifies the start of a new dictionary. The FLUSH codeword signifies that the DC chip has flushed out its buffer ie. it passes through the data currently held in the buffer without compressing that data prior to filling the buffer again with successive data and recommencing data compression. The DC
15 chip inserts RESET and FLUSH codewords into the data stream in an algorithm-dependent manner. However, the tape format places constraints on when certain RESET and FLUSH codewords must occur and also ensures the writing of certain information so as to enable the utilisation of certain ones
20 of the RESET and FLUSH codewords in order to improve access to the compressed data.

Decompression can only begin from a RESET codeword because the dictionary has to be rebuilt from the data. However, decompression can then stop at any subsequent FLUSH
25 codeword even though this is not at the end of that particular dictionary. This is why it is advantageous to put FLUSH codewords at the end of each record so as to enable selective decompression of segments of data which are smaller than that used to build a dictionary.

30 At the beginning of a dictionary, the majority of the data is passed through the DC chip without compression because most of the data will not have previously been seen. At this stage, the compression ratio is relatively small. Therefore, it is not desirable to have to restart a
35 dictionary so often as to reduce compression efficiency.

The main effect of putting extra information into the datastream is to reduce coupling between the data

compression engine and the system controller. Therefore, the only information which belongs in the datastream is that which is not directly needed by the controller, but is potentially of value to the decompression process.

5 Error checking information is perhaps the best example of information which can go into the datastream. It can get inserted during compression, and checked upon decompression. A CRC is a good example of this.

10 CRC stands for "Cyclic Redundancy Check." It is a syndrome generated by a series of bytes. It is used by some data transmission methods to provide a check that data corruption has not occurred during the transmission. It would be generated and sent immediately following the data. The receiver of the data would also generate it, and then
15 verify that its value matched the one received from the transmitter. If a four-byte CRC were used, for example, the chance of there being undetected errors would equal 2 to the 32nd.

20 The CRC is also used in data storage, where it is generated and written to the tape. The read process, then, generates its own and compares it with the one read from tape.

25 If a CRC were put into the datastream, it can only go after all the data that is used to generate it. Two choices still exist, however.

1. It can be compressed along with the data in the record.

2. It can be inserted into the datastream uncompressed after the record.

30 If the CRC is compressed, the value of it would not be available to non-decompressing drives. The tape format specification, for instance, would not need to leave room for it. It would be a part of the compression algorithm (rather than ancillary information in the data format).
35 However, because it is a function of uncompressed bytes, the CRC will be available to the hardware on the uncompressed side of the compression engine.

If the CRC is inserted, uncompressed, into the datastream, the CRC would be truly ancillary, and specified as such by the format definition. It would be available to any drive understanding the format, without any information
5 about the compression algorithm.

In the interest of compression efficiency, the second choice is better, since it is unlikely that the CRC would be found in ANY dictionary (- it is essentially a pseudo-random number which is a function of the bytes generating it). If
10 compressed, then, it will expand to approximately 1.5 times its original size. A four-byte CRC would take up to 6 bytes of storage if compressed, and only 4 if uncompressed. In the interest of reduced coupling, however, the first choice is better.

15 Another example of the type of information which might fit into the datastream is information which could be used to separate the data later. If the system controller does not need this information during writing or reading, and it can be generated on writing and skipped over on reading in
20 a simple fashion, it fits well into the datastream.

This sort of information is of value in a compressing environment, since even identically-sized records will produce variable-length records when compressed.

The Flush/EOR codeword is a data separator,
25 automatically inserted by the DC chip and removed by it. Only decompressing drives have access to these separations, however. Extra separation information would have to be included in the datastream for non-decompressing drives to have access to these boundaries. This would be ancillary
30 information.

This could be a compressed byte count (CBC). If the CBC were inserted into the format after the EOR codeword and uncompressed, a non-decompressing drive could use these as pointer information in a linked-list. Starting at the end
35 of a collection of compressed records, each having a CBC at the end, it would walk into the data and calculate where each compressed record in the collection begins and ends.

If both are used (EOR codewords/CBCs), the format is redundant. This redundancy can provide another check for the validity of the data that has been decompressed. The decompressor could compare the number of bytes it decompressed with the count in the datastream and signal an error if they did not match.

Methods for the storage of data, whether compressed or uncompressed, on tape will now be described.

The supply of the data from a user (host computer) to a tape storage apparatus will generally be accompanied by user separation of the data, whether this separation is the physical separation of the data into discrete packages (records) passed to the storage apparatus, or some higher level conceptual organisation of the records which is expressed to the storage apparatus by the host in terms of specific signals. This user-separation of data will have some particular significance to the host (though this significance will generally be unknown to the tape storage device). It is therefore appropriate to consider user separation as a logical segmentation even though its presence may be expressed to the storage apparatus through the physical separation of the incoming data.

Figure 1 (a) illustrates a sequence of user data and special separation signals that an existing type of host might supply to a tape storage apparatus. In this example, data is supplied in variable-length records R1 to R9; the logical significance of this physical separation is known to the host but not to the storage apparatus. In addition to the physical separation, user separation information is supplied in the form of special "file mark" signals FM. The file marks FM are provided to the storage apparatus between data records; again, the significance of this separation is unknown to the storage apparatus. The physical separation into records provides a first level of separation while the file marks provide a second level forming a hierarchy with the first level.

Figure 1 (b) shows one possible physical organisation for storing the user data and user separation information of Figure 1 (a) on a tape 10, this organisation being in accordance with a known data storage method. The mapping
5 between Figure 1 (a) and 1 (b) is straightforward - file marks FM are recorded as fixed-frequency bursts 1 but are otherwise treated as data records, with the records R1-R9 and the file marks FM being separated from each other by inter-block gaps 2 where no signal is recorded. The inter-
10 block gaps 2 effectively serve as first-level separation marks enabling the separation of the stored data into the user-understood logical unit of a record; the file marks FM (fixed frequency burst 1) form second-level separation marks dividing the records into logical collections of records.

15 Figure 1 (c) shows a second possible organisation which is known for storing the user data and user separation information of Figure 1 (a) on tape 10. In this case, the user data is organized into fixed-size groups 3 each including an index 4 for containing information about the
20 contents of the group. The boundary between two groups 3 may be indicated by a fixed frequency burst 5. The division of data into groups is purely for the convenience of the storage apparatus concerned and should be transparent to the host. The user data within a group is not physically
25 separated in any way and each record simply continues straight on from the end of the preceding one; all information regarding separation of the data in a group both into records and into the collection of records delimited by file marks is contained in the index of the group. In the
30 present example, records R1 to R8 and the first part of R9 are held in the illustrated group 3.

The length of the index 4 will generally vary according to the number of separation marks present and the number of records in the group; however, by recording the
35 index length in a predetermined location in the index with respect to the group ends, the boundary between the index and the last byte can be identified. A space with undefined

contents, eg. padding, may exist between the end of the data area and the first byte of the index.

The contents of the index 4 are shown in Figure 2 and, as can be seen, the index comprises two main data structures, namely a group information table 6 and a block access table 7. The number of entries in the block access table 7 is stored in a block access table entry (BAT ENTRY) count field in the group information table 6. The group information table 6 also contains various counts, such as a file mark count FMC (the number of file marks written since a beginning of recording (BOR) mark including any contained in the current group) and record counts RC (to be defined).

The block access table 7 describes by way of a series of access entries, the contents of a group and, in particular, the logical segmentation of the user data held in the group (that is, it holds entries indicative of each record boundary and separator mark in the group). The access entries proceed in order of the contents of the group.

Referring to Figure 3, the entries in the block access table each comprise a FLAG entry indicating the type of the entry and a COUNT entry indicating its value. The FLAG field is 8 bits and the COUNT field is 24 bits. The bits in the FLAG field have the following significance:

- | | | | |
|----|-----|---|--|
| 25 | SKP | - | A SKIP bit which, when set, indicates a "skip entry". A skip entry gives the number of bytes in the group which is not taken up by user data ie. the |
| 30 | | | size of the group minus the size of the user data area. |
| | XFR | - | A DATA TRANSFER bit which, when set, indicates the writing to tape of user data. |
| 35 | EOX | - | An END OF DATA TRANSFER bit which, when set, indicates the end of writing a user data record to tape. |

	CMP	-	A COMPRESSION bit which, when set, indicates that the entry relates to compressed data.
5	EOT	-	The value of this bit does not matter for the purposes of this description.
	MRK	-	A SEPARATOR MARK bit which, when set, indicates that the entry relates to a separator mark rather than to a data record.
10	BOR	-	A BEGINNING OF RECORD bit which, when set, indicates the location of the beginning of a data record.
15	EOR	-	An END OF RECORD bit which, when set, indicates the location of the end of a data record on tape.

Figure 3 illustrates the seven types of entry which can be made in the block access table. The SEPARATOR MARK entry has the BOR and EOR bit set because it is defined as a record. The next four entries each have the XFR bit set because they represent information about data transfers. The START PART OF RECORD entry relates to a case where only the beginning of a record fits into the group and the next part of the record runs over to the following group. The only bit set in the MIDDLE PART OF RECORD entry flag is the data transfer bit because there will not be a beginning or end of a record in that group. The END PART OF RECORD entry does not have the EOR bit set in the FLAG - instead, the EOR bit is set in the TOTAL COUNT entry which gives the total record byte count. The last entry in the block access table for a group is always a SKIP entry which gives the amount of space in the group which is not taken up by user data ie. the entry in the Count field for the SKIP entry equals the group size (eg. 126632 bytes) minus the data area size.

35 An example of a block access table for the group 3 of records shown in Figure 1 (c) is shown in Figure 4. The count entries for records R1-8 are the full byte counts for

those records whereas the count entry for record R9 is the byte count of the part of R9 which is in the group 3. The count entries for the file marks FM will be 0 or 1 according to the format. The count entry for the SKIP entry is 126632
5 minus the sum of the byte counts appearing previously in the table (not including Total Count entries).

In another embodiment there is a further possible entry in the block access table which signifies the algorithm used to compress the data in the group as shown in
10 Figure 3A. The algorithm number which is entered in the COUNT field is preferably one which conforms to a standard for DC algorithm numbers. The data transfer and total count FLAG entries for compressed records in the group have the CMP bit set. Thus compressed and uncompressed records in a
15 group can be distinguished by a drive on the basis of the CMP bit. For example, if we suppose that in Figure 1 (c), the even-numbered records are compressed records and the odd-numbered records are uncompressed, the block access table entries would be as shown in Figure 4A. In Figure 4A,
20 UBCX indicates an uncompressed byte count for record X and CBCX indicates a compressed byte count for record X.

Fig 5 shows another possible organisation for storing user data and related information on tape. Again, the user data is organised into fixed size groups each group
25 including an index (which is uncompressed even if the group contains compressed data) comprising a block access table for containing information about the contents of the group. The boundaries between groups may be indicated by fixed frequency bursts.

30 However, rather than storing information in the group index solely in terms of records, this embodiment involves storing the information about the contents of the group in terms of "Entities", where an entity comprises one or more records. In this embodiment, an entity can contain n
35 compressed records each having the same uncompressed length, where n is equal to or greater than 1.

In Figure 5, a group G comprises a single entity ENTITY 1 (or E_1) which comprises four complete records $CR_1 - CR_4$ of compressed data and a header portion H of 8 bytes. The records $CR_1 - CR_4$ have the same uncompressed length but may well be of different length after undergoing data compression.

The header portion H, which remains uncompressed, in the datastream contains the following information:

- 10 H_1 - The header length (4 bits). (The next 12 bits are reserved).
- ALG# - A recognised number denoting the compression algorithm being used to compress data (1 byte).
- 15 UBC - The uncompressed byte count for the records in the entity (3 bytes).
- #RECS - The number of records in the entity (2 bytes).

Optionally, an entity may include trailer portions at the end of each of the records in the entity, the trailer portions containing the compressed byte count of each record. Thus the trailer would occur immediately after an "end of record" (EOR) codeword. If this feature is present, the length of the trailer e.g. 3 bytes, could also be indicated in the header portion, in the 12 bits reserved after the header length H_1 .

An example of an embodiment in which each record in an entity has a trailer portion is shown in Figure 5A. The trailer portion is inserted into the datastream, uncompressed, at the end of each compressed record. Thus the entity in Figure 5A comprises a header portion H and four compressed records $CR_1 - CR_4$ of equal length when uncompressed, each of which has an uncompressed trailer portion T.

35 The trailer portion T of each record contains the compressed byte count (CBC) of the record and a cyclic redundancy check (CRC). The trailer occupies 6 bits at the

end of each record in this example. The length (T_1) of the trailer is included in the header portion H and occupies the last four bits of the first byte of the header portion H.

The inclusion of trailer portions does not alter the nature of the entries in the block access table 13 although the SKIP count entry will accordingly be smaller.

Insertion of compressed byte counts in the datastream has the advantage that a DC drive or a suitably configured non-DC drive can use these as pointers in a linked list to deduce where each compressed record begins and ends.

The use of both EOR codewords and CBC's in a DC-drive provides redundancy which can be utilised for error-checking purposes during decompression. The decompressor can signal an error if the CBC and the number of bytes which it decompressed do not match.

An advantage of including the length of the header portion (and the trailer portion if appropriate) in the header is that it enables this length to be varied whilst still allowing a drive to skip over the header if desired.

Information is recorded in a block access table T in the index of each group in terms of entities rather than in terms of records but otherwise as previously described with reference to Figures 2 - 4. The entries in the block access table for the entity E_i , are also shown in Figure 5.

The types of entries which are made in the block access table T are similar to those described with reference to Figure 2 - 4. The difference is that, now setting of the CMP bit in the FLAG field indicates that the entry relates to a byte count for an entity rather than for a record.

One possibility is to allow entities to contain only compressed records and this is preferred. This then means that the setting of the CMP bit in the FLAG field still indicates that the COUNT entry is a compressed byte count. However, another possibility is to allow entities to contain either compressed data or uncompressed data and to reserve a particular algorithm number eg. all zeros, to indicate that the data in an entity is uncompressed.

Storing information in the block access table T in terms of entities rather than records reduces the storage management overhead associated with writing and reading the records to and from tape. Whereas, using the scheme shown
5 in Figures 2 to 4, five entries in the block access table would be required for the group G, only two entries are now needed.

The organisation of records into entities facilitates the transfer of multiple records of identical uncompressed
10 size because it reduces the degree of processor intervention which is required during reading and writing. To write a sequence of records contained in an entity only requires processor intervention to form the header portion and to make the appropriate entry in the block access table. In
15 contrast, using the known scheme described with reference to Figures 1 to 4 requires processor intervention on a per record basis. This is especially important with data compression, since the compressed byte count is unknown until after the compression process has finished. Thus,
20 when trying to fill up a group with data, the number of records (and corresponding block access table entries) that will fit is unknown. By fixing the block access table requirements at one entry no matter how many records worth of data fit into the group, the entire group may be filled
25 up with a single processor intervention. Similar advantages are afforded when reading data.

With reference to Figure 6, an entity (E_n) may spread over more than one group eg. an entity E_1 containing a single, relatively long record CR_1 fills group G_1 and runs
30 over into group G_2 . The entries in the block access tables T_1 , T_2 of the groups G_1 , G_2 are also shown in Figure 6. To reduce the degree of linkage between groups, a new entity is started as soon as possible in a group ie. at the start of the group or at the beginning of the first compressed record
35 in the group if the previous record is uncompressed or at the beginning of the first new compressed record if the previous record is compressed and has run over from the

previous group. Therefore, at the end of compressed record CR_1 , the next entity, E_2 begins. Entity E_2 contains four compressed records CR_2 to CR_5 of equal uncompressed length.

It is envisaged that groups may contain a mixture of
5 entities containing compressed data and "naked records" containing uncompressed data. An example of this arrangement is shown in Figure 7 which also shows the corresponding entries in the block access table.

A group G contains an entity comprising a header
10 portion H and three compressed records CR_1 , CR_2 and CR_3 . The group G also comprises an uncompressed record R_4 (which has no header portion). The block access table T of the group G contains four entries:

the first entry is the full byte count of the entity
15 in the group;

the second entry is a file mark entry (which indicates the presence of a file mark in the incoming data before the start of record R_4);

the third entry is the full byte count of the
20 uncompressed record R_4 ;

the last entry is a SKIP entry.

It will be noted from Figure 7 that the CMP bit (the fourth bit of the FLAG field) is set for the entity byte count entry but not for the naked record byte count entry.

25 A suitably configured non-DC drive can identify compressed and uncompressed data on a tape having a mixture of such data by checking whether the CMP bit is set in the relevant block access table entries.

In this scheme, no separator marks are allowed within
30 an entity. For example, if a host is sending a sequence of equal length records to a DC tape drive and there is a file mark or other separator mark within that sequence, then the first set of records before the separator mark will be placed in one entity, the separator mark will be written to
35 tape and the set of records in the sequence which follow the file mark will be placed in a second entity. The corresponding entries for the two entities and the separator

mark will of course be made in the block access table of the relevant group (assuming that only one group is involved in this example)

The possible valid sequences of entries in the block access table of a group are illustrated in Figure 8. In Figure 8, states and actions are designated by rectangles and block access table entries are designated by ellipses. A 'spanned' record/entity is one which extends over from one group into another.

10 To take account of the existence of entities and the permitted existence of multiple compressed records within an entity, certain fields in the group information table in the index of each group are defined as follows:

Record Count - this field is a 4-byte field which specifies the sum of the values of the Number of Records in Current Group entry (see below) of the group information table of all groups up to and including the current group.

Number of Records in Current Group - this field is a 2-byte field which specifies the sum of the following:

20 i) the number of Separator Mark entries in the block access table of the current group.

ii) the number of Total Count of uncompressed record entries in the block access table of the current group.

25 iii) the number of Full Count uncompressed record entries in the block access table of the current group.

iv) the sum of the numbers of compressed records within all entities for which there is a Total Count of Entity entry or Full Count of Entity entry in the block access table of the current group.

30 v) the number, minus one, of compressed records in the entity for which there is a Start Part of Entity entry in the block access table of the current group, if such an entry exists.

vi) the number of Total Count of Entity entries in the block access table of the current group.

Group Number of the Previous Record - this field is a 2-byte field which specifies the running number of the

highest-numbered previous group in which a separator mark, an access point or the beginning of an uncompressed record occurred. It shall contain all ZERO bits if no such previous group exists.

5 With regard to the organisation of records in fixed size groups as described with reference to Figures 1 to 8 it is generally desirable to keep the groups independent from one another for decompression purposes ie. it is generally desirable to put a RESET codeword at or near the beginning
10 of each group. Two main reasons for this are to help reduce the amount of buffer space which is required in the controller by decreasing the linkages between groups ie. to make it less likely to have to store more than one group in the buffer at any one time. Another reason for putting a
15 RESET codeword at the beginning of a group is that, when it is desired selectively to decompress a record in the middle of a group it is not necessary to go outside the group to start the relevant dictionary.

 There are advantages in placing a FLUSH codeword after
20 each record - the FLUSH codeword is also called the "end of record" (EOR) codeword so as to improve the access to compressed data. This feature enables records to be decompressed individually, subject to the need to decompress from the RESET codeword which precedes the record. Having
25 a FLUSH codeword at the end of each record means that the data for each record can be decompressed without running into the data from the next record.

 The amount of compressed data which makes up a data dictionary is termed a "compression object". A compression
30 object may encompass more than one group of data as illustrated in Figure 9. Where a record overlaps from one group to the next, a RESET codeword is placed in the data stream at the beginning of the very next compressed record.

 In Figure 9 a Group G_1 comprises three full compressed
35 records CR_1 , CR_2 , CR_3 and the first part of a fourth compressed record CR_4 . The last part of record CR_4 extends

into the next group G_2 . The records are not organised into entities in this example.

During data compression, the dictionary is reset (indicated by R in Figure 9) at the beginning of group G_1 .

5 FLUSH codewords (indicated by F) are inserted into the datastream at the end of each record. The current dictionary continues until record CR_x ends at which time the dictionary is reset. Thus the current compression object comprises records $CR_1 - CR_x$.

10 If it is later desired selectively to decompress, say, record CR_3 , this can be achieved by beginning decompression at the start of record CR_1 ie. the start of the compression object containing record CR_3 , and decompressing data until the end of record CR_3 . A 'clean break' at the end of record
15 CR_3 can be achieved ie. without running over into the start of record CR_x due to the FLUSH codeword at the end of record CR_3 .

Thus, providing FLUSH codewords which are accessible by the format interspersed between 'access points' (RESET
20 codewords accessible by the format) enables selective decompression of segments of data which are smaller than the amount of data used to build a dictionary during data compression. The FLUSH codewords at the end of records are accessible since the compressed byte counts for each record
25 are stored in the block access table.

In the format, the start of a compression object which forms an 'access point' ie. a point at which the drive can start a decompression operation, may be denoted in one of several ways. Access points may be explicitly noted in the
30 block access table of each group. Alternatively, the presence of an access point may be implied by another entry in the block access table eg. the very presence of an algorithm number entry may imply an access point at the beginning of the first new record in that group.
35 Alternatively, a bit in the algorithm number may be reserved to indicate that a new dictionary starts at the beginning of the first new record in that group.

When records are organised into entities and entities are organised into groups as described with reference to Figures 5 to 7, a compression object may encompass more than one entity as illustrated in Figure 10, so as to obtain the advantage of dictionary sharing over entities which contain relatively small amounts of data.

Figure 10 shows three fixed size groups G_1 , G_2 , G_3 of compressed data. Group G_1 contains full record CR_1 and the first part of the next record CR_2 . Record CR_1 is the only record in entity E_1 . Group G_2 contains the middle part of record CR_2 . Group G_3 contains the end part of record CR_2 and contains further records CR_3 etc. Entity E_2 contains a single, relatively long record CR_2 .

During compression, the dictionary is reset (denoted by R) at the beginning of group G_1 but, since record CR_1 is relatively small, the compression object continues beyond record CR_1 and entity E_1 and includes record CR_2 and entity E_2 . A compression object ends at the end of record CR_2 and a new one begins at the beginning of record CR_3 .

A further possibility is for the presence of a non-zero algorithm number in an entity header to indicate the start of a new dictionary and otherwise for the algorithm number header entry to take a predetermined value eg. zero.

The presence of a FLUSH codeword at the end of each entity which is accessible owing to writing the compressed byte count of the entity in the block access table enables selective decompression of records on a per entity basis. For example, referring to Figure 10, the contents of entity E_2 (which happen to be a single record CR_2 in this example) could be decompressed without obtaining data from the beginning of record CR_3 . However, decompression must commence from the RESET codeword at the beginning of entity E_1 which is the nearest previous dictionary start point which is accessible in the tape format. It is also possible to decompress data on a per record basis utilising information in the entity header as will be described with reference to Figures 20A and 20B.

It should be appreciated that the DC chip inserts RESET codewords into the datastream in an algorithm-dependent manner - even in the middle of records. The above
5 description relates to the RESET codewords which are forced, recognised and utilised by the tape format.

To clarify, in Figures 5 to 10 the entities and compression objects do not include the indices of any relevant group.

10 A tape format for helical-scan implementation of the present invention will now be described.

The storage method and apparatus described hereinafter utilises a helical-scan technique for storing data in a format similar to that used for the storage of PCM audio
15 data according to the DAT Conference Standard (March 1988, Electronic Industries Association of Japan, Tokyo, Japan). The present method and apparatus is, however, adapted for storing computer data rather than digitised audio information.

20 Figure 11 shows the basic layout of a helical-scan tape deck 11 in which tape 10 from a tape cartridge 17 passes at a predetermined angle across a rotary head drum 12 with a wrap angle of 90° . In operation, the tape 10 is moved in the direction indicated by arrow T from a supply reel 13
25 to a take-up reel 14 by rotation of a capstan 15 against which the tape is pressed by a pinch roller 16; at the same time, the head drum is rotated in the sense indicated by arrow R. The head drum 12 houses two read/write heads HA, HB angularly spaced by 180° . In known manner, these heads
30 HA, HB are arranged to write overlapping oblique tracks 20, 21 respectively across the tape 10 as shown in Figure 12. The track written by head HA has a positive azimuth while that written by head HB has a negative azimuth. Each pair of positive and negative azimuth tracks, 20, 21 constitutes
35 a frame.

The basic format of each track as arranged to be written by the present apparatus is illustrated in Figure

12. Each track comprises two marginal areas 22, two sub areas 23, two ATF (Automatic Track Following) areas 24, and a main area 25. The ATF areas 24 provide signals enabling the heads HA, HB to accurately follow the tracks in known
5 manner. The main area 25 is used primarily to store the data provided to the apparatus (user data) although certain auxiliary information is also stored in this area; the sub areas 23 are primarily used to store further auxiliary information. The items of auxiliary information stored in
10 the main and sub areas are known as sub codes and relate for example, to the logical organisation of the user data, its mapping onto the tape, certain recording parameters (such as format identity, tape parameters etc), and tape usage history.

15 A more detailed description of the main area 25 and sub areas 23 will now be given including details as to block size that are compatible with the aforementioned DAT Conference Standard.

The data format of the main area 25 of a track is
20 illustrated in Figure 13. The main area is composed of 130 blocks each thirty six bytes long. The first two blocks 26 are pre-ambls which contain timing data patterns to facilitate timing synchronisation on playback. The remaining 128 blocks 27 make up the 'Main Data Area'. Each
25 block 27 of the Main Data Area comprises a four-byte 'Main ID' region 28 and a thirty-two byte 'Main Data' region 29, the compositions of which are shown in the lower part of Figure 13.

The main ID region 28 is composed of a sync byte, two
30 information-containing bytes W1, W2 and a parity byte. Byte W2 is used for storing information relating to the block as a whole (type and address) while byte W1 is used for storing sub codes.

The Main Data region 29 of each block 27 is composed
35 of thirty two bytes generally constituted by user-data and/or user-data parity. However, it is also possible to store sub codes in the Main Data region if desired.

The data format of each sub area 23 of a track is illustrated in Figure 14. the sub area is composed of eleven blocks each thirty-six bytes long. the first two blocks 30 are pre-ambles while the last block 31 is a post-
5 amble. The remaining eight blocks 32 make up the "Sub Data Area". Each block 32 comprises a four-byte 'Sub ID' region 33 and a thirty-two byte 'Sub Data' region 34, the compositions of which are shown in the lower part of Figure 14.

10 The Sub ID region 33 is composed of a sync byte, two information-containing bytes SW1, SW2 and a parity byte. Byte SW2 is used for storing information relating to the block as a whole (type and address) and the arrangement of the Sub Data region 34. Byte SW1 is used for storing sub
15 codes.

 The Sub Data region 34 of each block 32 is composed of thirty two bytes arranged into four eight-byte "packs" 35. These packs 35 are used for storing sub codes with the types of sub code stored being indicated by a pack-type label that
20 occupies the first half byte of each pack. The fourth pack 35 of every even block may be set to zero or is otherwise the same as the third pack while the fourth pack of every odd block is used to store parity check data for the first three packs both of that block and of the preceding block.

25 In summary, user data is stored in the Main Data regions 29 of the Main Data Area blocks 27 of each track while sub codes can be stored both in the Sub ID and Sub Data regions 33, 34 of Sub Data Area blocks 32 and in the Main ID and Main Data regions 28, 29 of Main Data Area
30 blocks 27.

 For the purposes of the present description, the sub codes of interest are an Area ID sub code used to identify the tape area to which particular tracks belong, and a number of sub codes used for storing counts of records and
35 separator marks. The area ID sub code is a four-bit code stored in three locations. Firstly, it is stored in the third and fourth packs 35 of the Sub Data region 34 of every

block in the Sub Data Areas of a track. Secondly, it is stored in byte SW1 of the Sub ID region 33 of every even Sub Data Area block 32 in a track, starting with the first block. The tape areas identified by this sub code will be
5 described later on with reference to Figure 15.

The sub codes used to store record and separator mark counts are stored in the first two packs 35 of the Sub Data region 34 of every block in the sub Data Areas of each track within the Data Area of the tape (see later with reference
10 to Figure 15). These counts are cumulative counts which are the same as the counts in the group information table as previously described. These counts are used for fast searching the tape and to facilitate this process are constant over a set of frames constituting a group, the
15 counts recorded in the tracks of a group of frames being the counts applicable as of the end of the group.

The general organisation of frames along the tape as implemented by the present storage method and apparatus will be considered next. Thus, referring to Figure 15, the tape
20 can be seen to be organised into three main areas, namely a lead-in area 36, a data area 37 and an end-of-data (EOD) area 38. The ends of the tape are referenced BOM (beginning of media) and EOM (end of media). User data is recorded in the frames of data area 37. The lead-in area 36 includes an
25 area between a beginning-of-recording BOR mark and the data area 37 where system information is stored. The Area ID sub code enables the system area, data area 37 and EOD area 38 to be distinguished from one another.

The frames 48 of the data area are arranged in groups
30 39 each of a fixed number of frames (for example, twenty two); optionally, these groups are separated from each other by one or more amble frames of predetermined content. In terms of organisation of user data records, these groups 39 correspond to the group 3 described with reference to Figure
35 1(c). Thus, the placement of user data into such groups 39 has no relation to the logical segmentation of the user data and information relating to this segmentation (record marks,

separator marks) is stored in an index 40 that terminates the user-data in a group (the index actually occupies user data space within the group). Note that although the index is shown in Figure 15 as occupying the final portion of the last frame of the group, this is only correct in relation to the arrangement of data prior to a byte-interleaving operation that is normally effected before data is recorded on tape; however, for present purposes, the interleaving operation can be disregarded.

10 In practice, the information in the index is physically dispersed within the main data areas of the tracks in the group.

The contents of the index 4 are shown in Figure 2 and, as previously described, the index comprises two main data structures, namely a group information table and a block access table. The group information table is stored in a fixed location at the end of the group and is the same size independent of the contents of the group. In contrast, the block access table varies in size depending on the contents of the group and extends from the group information table backwards into the remainder of the user data area of the frames of the group. Entries are made in the block access table from the group information table backwards to the boundary with real user data or 'pad'.

25 Also shown in Figure 15 are the contents of a sub data area block 32 of a track within a data-area group 39. As previously noted, the first two packs contain a separator mark count, the second pack 35 also contains record counts RC (as defined above), and the third pack 35 contains the Area ID and an absolute frame count AFC. For all the tracks in a group, the counts FMC, and RC held in the sub data area blocks are the same as those held in the group information table 41 of the group index 40.

35 Figure 16 is a block diagram of the storage apparatus for compressing and recording user data in accordance with the above-described tape format. The apparatus includes the tape deck 11 already described in part with reference to

Figure 11. In addition to the tape deck, the apparatus includes an interface unit 50 for interfacing the apparatus with a host computer (not shown) via a bus 55; a group processor 51 comprising a data compression processor (DCP) and a frame data processor 52 for processing user-record data and separation data into and out of Main Data Area and Sub Data Area blocks 27 and 32; a signal organiser 53 for composing/decomposing the signals for writing/reading a track and for appropriately switching the two heads HA, HB; and a system controller 54 for controlling the operation of the apparatus in response to commands received from a computer via the interface unit 50. Each of the main component units of the apparatus will be further described below.

15 Firstly, the structure and operation of the data compression processor (DCP) or data compression engine will be described.

With reference to Figure 17 the heart of the engine is a VLSI data compression chip (DC chip) which can perform both compression and decompression on the data presented to it. However, only one of the two processes (compression or decompression) can be performed at any one time. Two first-in, first-out (FIFO) memories are located at the input and the output of the DC chip to smooth out the rate of data flow through the chip. The data rate through the chip is not constant, since some data patterns will take more clock cycles per byte to process than other patterns. The instantaneous data rate depends upon the current compression ratio and the frequency of dictionary entry collisions, both of which are dependent upon the current data and the entire sequence of data since the last dictionary RESET. The third section of the subsystem is a bank of static RAM forming an external dictionary memory (EDM) that is used for local storage of the current dictionary entries. These entries contain characters, codeword pointers, and control flags.

35 Fig 18 shows a block diagram of the DC integrated circuit. The DC chip is divided into three blocks; the

input/output converter (IOC), the compression and decompression converter (CDC), and the microprocessor interface (MPI).

The MPI section provides facilities for controlling and observing the DC chip. It contains six control registers, eight status registers, two 20 bit input and output byte counters, and a programmable automatic dictionary reset circuit. The control and status registers are accessed through a general-purpose 8 bit microprocessor interface bus. The control registers are used to enable and disable various chip features and to place the chip into different operating modes (compression, decompression, pass through, or monitor). The status registers access the 20 bit counters and various status flags within the chip.

It has been found that compression ratios can be improved by resetting the dictionary fairly frequently. This is especially true if the data stream being compressed contains very few similar byte strings. Frequent dictionary resets provide two important advantages. First, resetting the dictionary forces the codeword length to return to 9 bits. Second, new dictionary entries can be made that reflect the present stream of data (a form of adaption). The DC chip's interface section contains circuitry that dynamically monitors the compression ratio and automatically resets the dictionary when appropriate. Most data compression algorithms will expand their output if there is little or no redundancy in the data.

The IOC section manages the process of converting between a byte stream and a stream of variable-length codewords (ranging from 9 bits to 12 bits). Two of the eight reserved codewords are used exclusively by the IOC. One of these codewords is used to tell the IOC that the length of the codewords must be incremented by one. Thus, the process of incrementing codeword size is decoupled from the CDC section - the IOC operates as an independent pipeline process, thus allowing the CDC to perform

compression or decompression without being slowed down by the IOC.

The second reserved codeword which is the FLUSH (or 'end of record' (EOR)) codeword alerts the IOC that the next codeword is the last one associated with the current packet of data ie. the FLUSH codeword is actually the penultimate one of a compressed record. From this information, the IOC knows to finish its packing routine and end on a byte boundary. This feature allows compression of multiple input packets into one contiguous output packet while maintaining the ability to decompress this packet into its constituent packets. The IOC is also capable of allowing data to pass straight through from input to output without altering it, and of allowing data to pass through while monitoring the potential compression ratio of the data. These features can be used as another level of expansion protection.

The CDC section is the engine that performs the transformation from uncompressed data to compressed data and vice versa. This section is composed of control, data path, and memory elements that are adjusted for maximum data throughput. The CDC interfaces with the IOC via two 12 bit buses. During compression, the IOC passes the input bytes to the CDC section, where they are transformed into codewords. These codewords are sent to the IOC where they are packed into bytes and sent out of the chip. Conversely, during decompression the IOC converts the input byte stream into a stream of codewords, then passes these codewords to the CDC section, where they are transformed into a stream of bytes and sent to the IOC. The CDC section also interfaces directly to the external RAM that is used to store the dictionary entries.

The CDC makes use of two reserved codewords. The first is used any time a dictionary reset has taken place. The occurrence of this codeword causes two actions: the IOC returns to the state in which it packs or unpacks 9 bit codewords, and the CDC resets the current dictionary and

starts to build a new one. Dictionary resets are requested by the MPI section via microprocessor control or the automatic reset circuitry. The second reserved codeword is generated during compression any time the CDC runs out of usable external RAM while trying to build a new dictionary entry. This event very rarely happens, given sufficient external RAM. However, as the amount of memory decreases, it is more likely that the CDC will encounter too many dictionary collisions and will not be able to build new dictionary entries. With the reduction of external memory and the inevitable increase in dictionary collisions, the data throughput and compression performance will be slightly degraded. This "full dictionary" codeword is also used during decompression by the CDC to ensure that the decompression process stops building dictionary entries at the same point as the compression process.

Returning now to Figure 16 the data storage apparatus is arranged to respond to commands from a computer to load/unload a tape, to store a data record or separation mark, to enable compression of data, to search for selected separation marks or records, and to read back the next record.

The interface unit 50 is arranged to receive the commands from the computer and to manage the transfer of data records and separation marks between the apparatus and computer. Upon receiving a command from the computer, the unit 50 passes it on to the system controller 54 which, in due course, will send a response back to the computer via the unit 50 indicating compliance or otherwise with the original command. Once the apparatus has been set up by the system controller 54 in response to a command from the computer to store or read data, then the interface unit 50 will also control the passage of records and separation marks between the computer and group processor 51.

During data storage the group processor 51 is arranged to compress the user-data if required and to organise the user-data that is provided to it in the form of data

records, into data packages each corresponding to a group of data. The processor 51 is also arranged to construct the index for each group and the corresponding sub codes. During reading, the group processor effects a reverse
5 process enabling data records and separation marks to be recovered from a group read from tape prior to decompression.

The form of the group processor 51 is shown in Figure 19. At the heart of the group processor 51 is a buffer 56
10 which is arranged to hold more than one (for example, two) group's worth of data. The allocation of buffer space to incoming and outgoing data is controlled by a buffer space manager 57. The processor 51 communicates with the interface 50 via a first interface manager 58 and with the
15 frame data processor 52 via a second interface manager 59. Overall control of the grouping process is effected by a grouping manager 60 which also generates the group indices and associated codes during recording (functional block 61) and interprets these indices and sub codes during reading
20 (functional block 62). The grouping manager 60 is arranged to exchange coordination signals with the system controller 54.

The DC processor DCP is operable to compress data for storage on tape or to decompress data to be read by a host.
25 There are interconnections between the DC processor DCP and the interface manager 58, the buffer 56, the buffer space manager 57 and the grouping manager 60 for the interchange of control signals.

The grouping manager 60 also comprises an entity
30 manager (EM) which organises compressed data into entities and generates header portions for the entities. The grouping manager 60 and the buffer space manager 57 are control components and data for writing to tape does not pass through them, but rather passes directly from the
35 buffer 56 to the interface manager 59.

During recording when the host is ready to pass a data record, the interface 50 asks the buffer space manager 57 (via the interface manager 58) whether the processor 51 is ready to receive the record. The buffer space manager 57 may initially send a 'wait' reply but, in due course, enables the transfer of the data record from the host to the buffer 56.

If the data is to be compressed (according to control signals from the system controller 54), the DC processor DCP substitutes codewords for a proportion of the data in the record in accordance with a data compression algorithm as previously described.

Typically, a host transfers records one at a time although multiple record transfers make sense for shorter records.

The grouping manager 60 is connected to the buffer space manager 57 and tells the buffer space manager 57 how much more data the group can take before it runs into the index area of the group. The buffer space manager 57 notifies the grouping manager 60 whenever the maximum number of bytes has been transferred into the current group or the last byte from the host has been received.

If a transfer from the host cannot all fit inside a group, it is said to "span" the group boundary. The first part of the transfer goes into one group and the rest into subsequent groups. The buffer space manager 57 tells the grouping manager 60 if the host tries to supply more data than will fit in the current group being built. If no span occurs, the group index is updated and the grouping manager 60 waits for another write command. If a span occurs, the index of the current group is updated and that group is available for writing to tape. The next group is begun and the data from the host goes directly into the beginning of that new group.

The record will be transferred to a buffer location that corresponds to the eventual positioning of the record data within the group of which it is to form a part.

Information on the size of the record is passed to the grouping manager 60. When the host sends a separator indication this is also routed to the grouping manager 60. The grouping manager keeps track of the separator mark and
5 record counts from BOR and uses this information in the construction of the index and separation-count and record count sub codes of a group. The index is constructed in a location in the buffer appropriate to its position at the end of a group.

10 In parallel, the entity manager EM generates an entity header portion for the current entity which will contain the compressed record data. The header portion is not compressed. Likewise, the entity manager EM may generate trailer portions (also uncompressed) for each record.

15 The entity manager EM is responsible for ensuring that the rules governing entity formation are observed. These are:-

- a) Start a new entity:
 - 20 i) as soon as possible after the beginning of a group;
 - ii) when the uncompressed size of records being sent from the host changes;
 - iii) when the compression algorithm changes, and
(Regarding i) and iii) above, the need for an access
25 point requires starting a new entity and an appropriate signal is sent to the data compression processor DCP from the grouping manager 60.)
- b) End an entity:
 - 30 i) when an uncompressed record requires to be stored;
 - ii) when a separation mark requires to be stored.

The formation of each entity triggers a BAT entry.

When a group becomes full, the processes of data
35 compression and entity building halt until a new group is initiated.

If incoming data is not to be compressed, the data passes unchanged through the DC processor DCP and the entity manager EM is inactive. Uncompressed records are organised directly into groups without forming part of an entity and information regarding the records is put into the group index. Uncompressed records do not have a header portion created for them.

Once a group (including its index and sub codes) has been assembled, it is transferred to the frame data processor 52 for organisation into the blocks making up the main data areas and sub data areas of twenty two successive frames. Information about frame ID is in the datastream. There is a continuous stream of data between the group processor 52 to a small buffer in the frame data processor 52 which is able to store three frame's worth of data.

As previously mentioned, it may be desirable to insert one or more amble frames between groups of frames recorded on the tape. This can be done by arranging for the frame data processor 52 to generate such amble frames either upon instruction from the group processor 51 or automatically at the end of a group if the processor 52 is aware of group structure.

By sizing the buffer 56 such that it can hold two group's worth of data, the general operation of the processor 51 can be kept as straight forward as possible with one group being read in and one group being processed and output. During writing, one group is being built with data from a host and one is being written to tape.

When data is being read from tape, the group processor 51 is arranged to receive user-data and sub-codes on a frame-by-frame basis from the frame data processor 52, the data being written into the buffer 56 in such a manner as to build up a group. The group processor 51 can then access the group index to recover information on the logical organisation (record/entity structure, separator marks) of the user-data in the group and an indication of whether the data is compressed.

If the data is uncompressed, or the data is compressed but is to be read back to the host in its compressed form for software decompression, the group processor 51 can pass a requested record or separator mark to the host via the
5 interface 50 in which case the data passes through the DC processor DCP unchanged. The entity header portions in compressed data are passed back to a host by a non-DC drive for use by the host.

If the data is compressed and is to be decompressed,
10 the data is decompressed by the DC processor DCP in the manner described above before being passed to the host.

The header portions from each entity are utilised by a DC drive but are not passed to the DC processor DCP. The algorithm number in the header portion is checked for
15 consistency with the algorithm used by the DC processor DCP. Further, the number of compressed records in the entity is obtained from the header portion enabling a record count down to be performed as the entity data is passed to the DC processor DCP.

20 To facilitate the assembly of frame data back into a group's worth of data, each frame can be tagged with an in-group sequence number when the frame is written to tape. This in-group number can be provided as a sub code that, for example, is included at the head of the main data region of
25 the first block in the Main Data Area of each track of a frame. The subcode is used on reading to determine where the related frame data is placed in the buffer 56 when passed to the group processor 51.

The frame data processor 52 functionally comprises a
30 Main-Data-Area (MDA) processor 65, a Sub-Data-Area (SDA) processor 66, and a sub code unit 67 (in practice, these functional elements may be constituted by a single microprocessor running appropriate processes).

The sub code unit 67 is arranged to provide subcodes
35 to the processors 65 and 66 as required during writing and to receive and distribute sub codes from the processors 65, 66 during reading. Depending on their information contents,

sub codes may be generated/required by the group processor 51 or the system controller 54; the separation mark count sub codes are, for example, determined/used by the group processor 51 while the Area ID sub codes are determined/used
5 by the controller 54. In the case of non-varying sub codes such as certain writing parameters, the sub codes may be permanently stored in the unit 67. Furthermore, any frame-dependent sub codes may conveniently be generated by the sub code unit 67 itself.

10 The MDA processor 65 is arranged to process a frame's worth of user data at a time together with any relevant sub codes. Thus during recording, the processor 65 receives a frame's worth of user-data from the group processor 51 together with sub codes from the unit 67. On receiving the
15 user-data the processor 65 interleaves the data, and calculates error correcting codes, before assembling the resultant data and sub codes to output the Main-Data-Area blocks for the two tracks making up a frame. In fact before assembling the user data with the sub codes, scrambling
20 (randomising) of the data may be effected to ensure a consistent RF envelope independent of the data contents of a track signal.

During reading, the processor 65 effects a reverse process on the two sets of Main-Data-Area blocks associated
25 with the same frame. Unscrambled, error-corrected and de-interleaved user data is passed to the group processor 51 and sub codes are separated off and distributed by the unit 67 to the processor 51 or system controller 54 as required.

The operation of the SDA processor 66 is similar to
30 the processor 65 except that it operates on the sub codes associated with the sub-data-areas of a track, composing and decomposing these sub codes into the from Sub-Data-Area blocks.

The signal organiser 53 comprises a
35 formatter/separator unit 70 which during recording (data writing) is arranged to assemble Main-Data-Area blocks and Sub-Data-Area blocks provided by the frame data processor 52

together with ATF signals from an ATF circuit 80, to form the signal to be recorded on each successive track. The necessary pre-amble and post-amble patterns are also inserted into the track signals where necessary by the unit

5 70. Timing signals for coordinating the operation of the unit 70 with rotation of the heads HA, HB are provided by a timing generator 71 fed with the output of a pulse generator 81 responsive to head drum rotation. The track signals output on line 72 from the unit 70 are passed alternately to

10 head HA and head HB via a head switch 73, respective head drive amplifiers 74, and record/playback switches 75 set to their record positions. The head switch 73 is operated by appropriately timed signals from the timing generator 71.

During playback (data reading) the track signals

15 alternately generated by the heads HA and HB are fed via the record/playback switches 75 (now set in their playback positions), respective read amplifiers 76, a second head switch 77, and a clock recovery circuit 78 to the input of the formatter/separator unit 70. The operation of the head

20 switch 77 is controlled in the same manner as that of the head switch 73. The unit 70 now serves to separate off the ATF signals and feed them to the circuit 80, and to pass the Main-Data-Area blocks and Sub-Data-Area blocks to the frame data processor 52. Clock signals are also passed to the

25 processor 52 from the clock recovery circuit 78.

The switches 75 are controlled by the system controller 54.

The tape deck 11 comprises four servos, namely a capstan servo 82 for controlling the rotation of the capstan

30 15, first and second reel servos 83, 84 for controlling rotation of the reels 14, 15 respectively, and a drum servo 85 for controlling the rotation of the head drum 12. Each servo includes a motor M and a rotation detector D both coupled to the element controlled by the servo. Associated

35 with the reel servos 83, 84 are means 86 for sensing the beginning-of-media (BOM) and end-of media (EOM); these means 86 may for example be based on motor current sensing, as the

motor current of whichever reel is being driven to wind in tape (dependent on the direction of tape travel) will increase significantly upon stalling of the motor at BOM/EOM.

5 The tape deck 11 further comprises the automatic track following circuit 80 for generating ATF signals for recordal on tape during recording of data. During reading, the ATF circuit 80 is responsive to the ATF track signal read from tape to provide an adjustment signal to the capstan servo 82
10 such that the heads HA, HB are properly aligned with the tracks recorded on the tape. The tape deck 11 also includes the pulse generator 81 for generating timing pulses synchronised to the rotation of the heads HA, HB.

 The operation of the tape deck 11 is controlled by a
15 deck controller 87 which is connected to the servos 82 to 85 and to the BOM/EOM sensing means 86. The controller 87 is operable to cause the servos to advance the tape, (either at normal speed or at high speed) through any required distance. This control is effected either by energising the
20 servos for a time interval appropriate to the tape speed set, or by feedback of tape displacement information from one or more of the rotation detectors D associated with the servos.

 The deck controller 87 is itself governed by control
25 signals issued by the system controller 54. The deck controller 87 is arranged to output to the controller 54 signals indicative of BOM and EOM being reached.

 The system controller 54 serves both to manage high-level interaction between the computer and storage apparatus and to coordinate the functioning of the other units of the
30 storage apparatus in carrying out the basic operations of Load/Write/Compress/Decompress/Search/Read/Unload requested by the computer. In this latter respect, the controller 54 serves to coordinate the operation of the deck 11 with the
35 data processing portion of the apparatus.

 In controlling the tape deck 11, the system controller can request the deck controller 87 to move the tape at the

normal read/write speed (Normal) or to move the tape forwards or backwards at high speed, that is Fast Forward (F.FWD) or Fast Rewind (F.RWD). The deck controller 87 is arranged to report arrival of BOM or EOM back to the system controller 54.

An operation to locate a record for decompression will now be described with reference to Figures 20A and 20B.

Upon the host issuing a command to decompress a record, the controller 54 generates a search key having a value equal to the record count of the record to be decompressed. The current record count is held in the grouping manager 60 of the group processor 51. Next the tape is advanced (or rewound as appropriate) at high speed (many times faster than normal) while the head drum is rotated at a speed to maintain the relative velocity of the heads HA, HB across the tape at a constant value; in this mode, it is possible to read the sub area of about one track in every three hundred (steps 91a and 91b). Reading track sub areas at speed is a known technique and will therefore not be described in detail.

Fast forward searching is depicted in Figure 20A and fast backward searching is depicted in Figure 20B.

During fast forward searching (Figure 20A), for each sub area that is successively read, the record count held in the second pack of each sub data area block is compared by the controller 54 with the search key (step 92a). If the record count is less than the search key, the search is continued; however, if the record count is equal to, or greater than the search key, fast forward searching is terminated and the tape is backspaced through a distance substantially equal to the distance between fast forward reads (step 93). This ensures that the record count held in the sub areas of the track now opposite the head drum will be less than the search key.

During fast backward searching (Figure 20B), for each sub area that is successively read, the record count held in the second pack of each sub data block is compared by the

controller 54 with the search key (step 92b). If the record count is more than the search key, the search is continued; however, if the record count is equal to or less than the search key, the fast rewind is stopped.

5 Next, for both fast forward and fast backward searching, the tape is advanced at its normal reading speed (step 94), and each successive group is read off tape in turn and temporarily stored in the buffer 56 of the group processor 51. The record count held in the index of each
10 group is compared with the search key (step 95) until the count first equals or exceeds the search key. At this point, reading is stopped as the record searched for is present in the group in buffer 56 whose record count has just been tested. If entries are made in the block access
15 table on a per record basis the block access table of the index of this group is now examined to identify the record of interest (step 96) and the address in the buffer of the first data record byte is calculated (step 97). Thereafter, the group processor 51 tells the system controller 54 that
20 it has found the searched-for record and is ready to decompress and read the next data record; this is reported back to the host by the controller (step 98). The search operation is now terminated.

It will, of course, be appreciated that other search
25 methods could be implemented.

In order to detect when the bounds of the data area of the tape have been exceeded while searching at speed, whenever a sub area is read the Area ID sub code is checked by the system controller 54. If this sub code indicated
30 that the searching has gone beyond the data area of the tape, then the tape direction is reversed and searching is resumed, generally at a lower speed. For clarity, this Area ID check has been omitted from Figures 20A and 20B.

The next step after the record of interest has been
35 located is to check the algorithm number indicating which algorithm was used to compress the data in the record. This

is done by examining the block access table of the relevant group if the algorithm number is stored in that table.

If the algorithm number corresponds to the algorithm used by the DC chip in the tape drive (or to one of the DC chips if there is more than one), the next step is to locate the beginning of the compression object containing the record of interest. This may be done in a variety of ways depending on the particular recording format as described with reference to Figure 9.

Once the beginning of the compression object containing the record of interest is found, decompression commences from that point and continues until the FLUSH (or EOR) codeword at the end of the record is reached. The decompressed record can then be passed to the host. The presence of a FLUSH codeword at the end of the record means that the record can be decompressed cleanly without obtaining data from the beginning of the next record.

If compressed records are organised into entities, the group of interest is located as described earlier with reference to Figures 20A and 20B.

The relevant entity can then be located by using the #RECS entries in the entity headers within the group. Decompression is started from the nearest previous access point which may be found by checking the algorithm ID entry in the relevant entity and, if it indicates that the compressed data in that entity is a continuation of an earlier started dictionary, skipping back to the previous entity header and so on until an access point is found. Only decompressed data obtained from the relevant record or records is retained. The existence of data in the entity headers therefore has the advantage of facilitating finding relevant records and access points and allows the process of data management to be decoupled from that of decompression. If there are trailers provided after each compressed record in an entity which contain the compressed byte count of the record, these CBCs can be utilised to advantage in ascertaining when to start retaining decompressed data

rather than (or as well as) counting FLUSH codewords during decompression.

Consequently, the presence of ancillary information in the data stream can be used to advantage in finding selected records, the nearest previous access point and in 5
ascertaining the point at which decompressed data should be kept.

During normal reading of data, the ancillary information eg. the error checking information and/or data 10
separation information in the datastream, is utilised accordingly. One possibility is for the drive (DC or non-DC) to generate CRCs and compare these with the CRCs in the trailer portions of records organised into entities. Also, the drive (again DC or non-DC) can use the CBCs in the 15
trailer portions to find out where each compressed record begins and ends.

It should be appreciated that the present invention is not limited to helical-scan data recording. The compression algorithm described is purely an example and the present 20
invention may also be applicable to the storage of data which is compressed according to a different algorithm.

25

30

35

CLAIMS

1. A data storage method for writing compressed data organised in the form of records (CR_n) to tape (10) characterised by inserting into the datastream ancillary information which is extra to the data compression process.
2. A method according to claim 1 wherein the ancillary information comprises error checking information.
3. A method according to claim 1 or 2 wherein the ancillary information comprises data separation information.
4. A method according to any preceding claim comprising writing the ancillary information to tape (10) in uncompressed form.
5. A method according to any preceding claim comprising inserting into the datastream ancillary information in association with one or more records.
6. A method according to claim 5 comprising inserting into the datastream a header portion (H) containing ancillary information relating to one or more records (CR_n) following the header portion.
7. A method according to claim 5 or claim 6 comprising inserting into the datastream a trailer portion (T) containing ancillary information relating to one or more records (CR_n) preceding the trailer portion.
8. A method according to any preceding claim comprising organising data records into groups (G_n) independently of the record structure of the data and writing information regarding the records (CR_n) in a group to an index (4) associated with the group.

9. A method according to claim 8 comprising writing information to the group indices in terms of entities (E_n), where an entity comprises one or more records (CR_n).

10. A method according to claim 9 comprising writing ancillary information to a header (H) associated with each entity (E_n).

11. A storage device for compressing user data and writing compressed data to tape (10) which is operable in accordance with a method as claimed in any preceding claim.




FIG A




FIG B

FIG 1






FIG 2



	FLAG										COUNT
	S	X	E	C	E	M	B	O	R	E	
SEPARATOR MARK	0	0	0	0	X	1	1	1	1	1	0,1 (SIGNIFY DIFFERENT MARK TYPES)
TOTALLY CONTAINED RECORD	0	1	1	0	X	0	1	1	1	1	FULL RECORD BYTE COUNT
START OF PART RECORD	0	1	0	0	X	0	1	0	1	0	FIRST PART BYTE COUNT
MIDDLE PART OF RECORD	0	1	0	0	X	0	0	0	0	0	MIDDLE PART BYTE COUNT
END PART OF RECORD	0	1	1	0	X	0	0	0	0	0	END PART BYTE COUNT
TOTAL COUNT	0	0	0	0	X	0	0	0	1	1	TOTAL RECORD BYTE COUNT
SKIP ENTRY	1	0	0	0	X	0	0	0	0	0	GROUP SIZE MINUS DATA AREA SIZE

FIG 3

	FLAG								COUNT	
	S	X	E	C	E	M	B	E		
	K	F	O	M	O	R	O	R		
	P	R	X	P	T	K	R	R		
ALGORITHM NUMBER	0	0	0	1	X	0	0	0	0	DC ALGORITHM #
SEPARATOR MARK	0	0	0	0	X	1	1	1	1	0,1 (SIGNIFY DIFFERENT MARK TYPES)
TOTALLY CONTAINED RECORD	0	1	1	1	X	0	1	1	1	FULL RECORD COMPRESSED BYTE COUNT
START OF PART RECORD	0	1	0	1	X	0	1	0	0	FIRST PART COMPRESSED BYTE COUNT
MIDDLE PART OF RECORD	0	1	0	1	X	0	0	0	0	MIDDLE PART COMPRESSED BYTE COUNT
END PART OF RECORD	0	1	1	1	X	0	0	0	0	END PART COMPRESSED BYTE COUNT
TOTAL COUNT	0	0	0	1	X	0	0	1	1	TOTAL RECORD COMPRESSED BYTE COUNT
SKIP ENTRY	1	0	0	0	X	0	0	0	0	GROUP SIZE MINUS DATA AREA SIZE

FIG 3A

7/24

FLAG	COUNT
0 1 1 0 X 0 1 1	BC 1
0 1 1 0 X 0 1 1	BC 2
0 0 0 0 X 1 1 1	FM
0 1 1 0 X 0 1 1	BC 3
0 1 1 0 X 0 1 1	BC 4
0 1 1 0 X 0 1 1	BC 5
0 0 0 0 X 1 1 1	FM
0 1 1 0 X 0 1 1	BC 6
0 1 1 0 X 0 1 1	BC 7
0 0 0 0 X 1 1 1	FM
0 1 1 0 X 0 1 1	BC 8
0 0 0 0 X 1 1 1	FM
0 1 0 0 X 0 1 0	BC 9a
1 0 0 0 X 0 0 0	SKIP COUNT

FIG 4

8 / 24

FLAG	COUNT
0 0 0 1 X 0 0 0	ALG #
0 1 1 0 X 0 1 1	UBC 1
0 1 1 1 X 0 1 1	CBC 2
0 0 0 0 X 1 1 1	FM
0 1 1 0 X 0 1 1	UBC 3
0 1 1 1 X 0 1 1	CBC 4
0 1 1 0 X 0 1 1	UBC 5
0 0 0 0 X 1 1 1	FM
0 1 1 1 X 0 1 1	CBC 6
0 1 1 0 X 0 1 1	UBC 7
0 0 0 0 X 1 1 1	FM
0 1 1 1 X 0 1 1	CBC 8
0 0 0 0 X 1 1 1	FM
0 1 0 0 X 0 1 0	UBC 9a
1 0 0 0 X 0 0 0	SKIP ENTRY

FIG 4A




FIG 5




FIG 5A




FIG 6


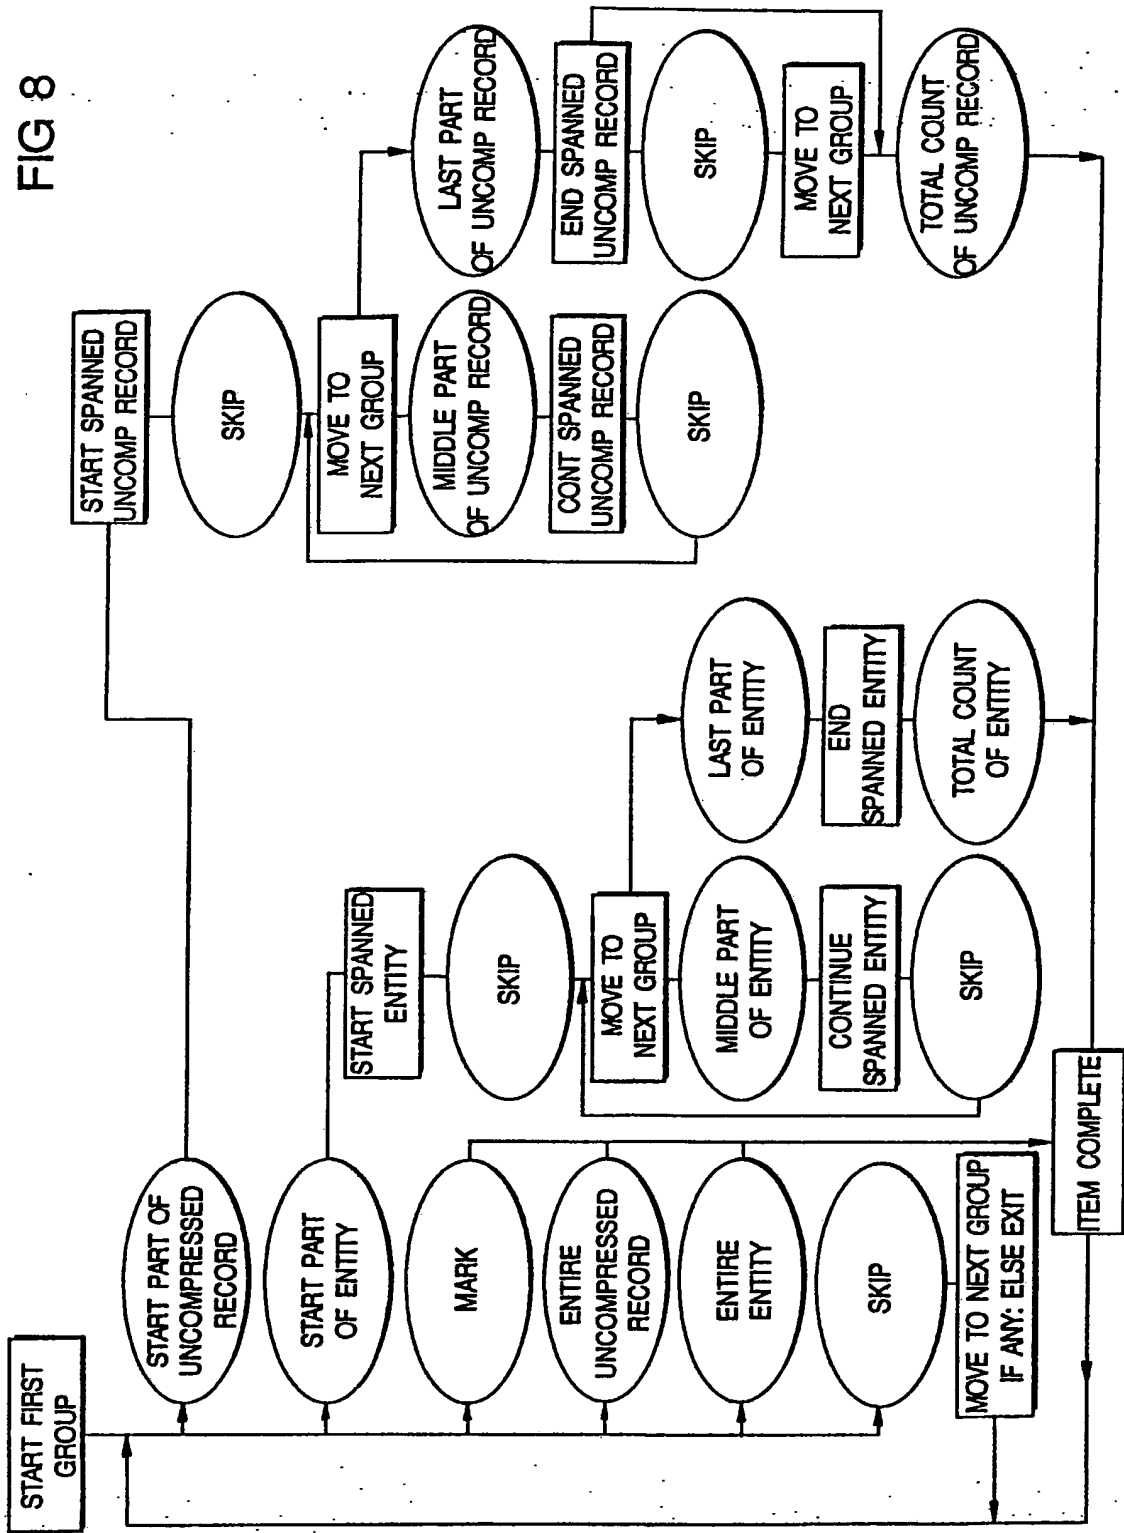


FIG 7

FIG 8






FIG 9




FIG 10

16/24




FIG 11






FIG 12

SUBSTITUTE SHEET






FIG 15




FIG 16

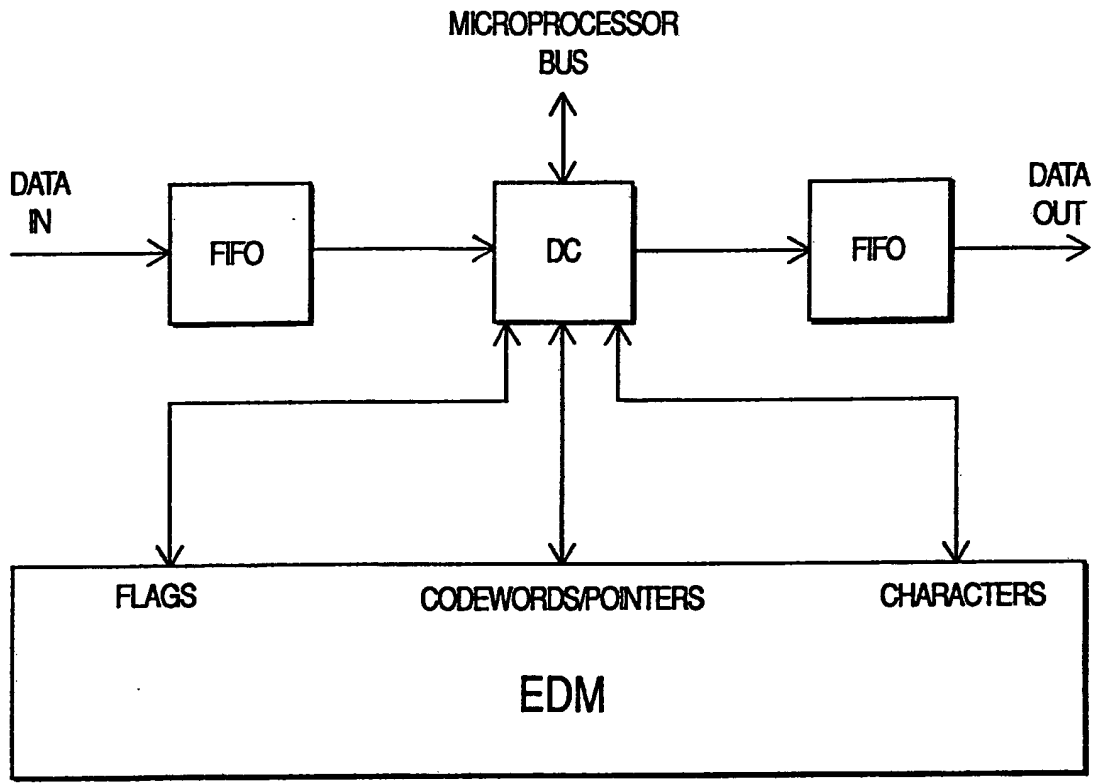


FIG 17




FIG 18




FIG 19

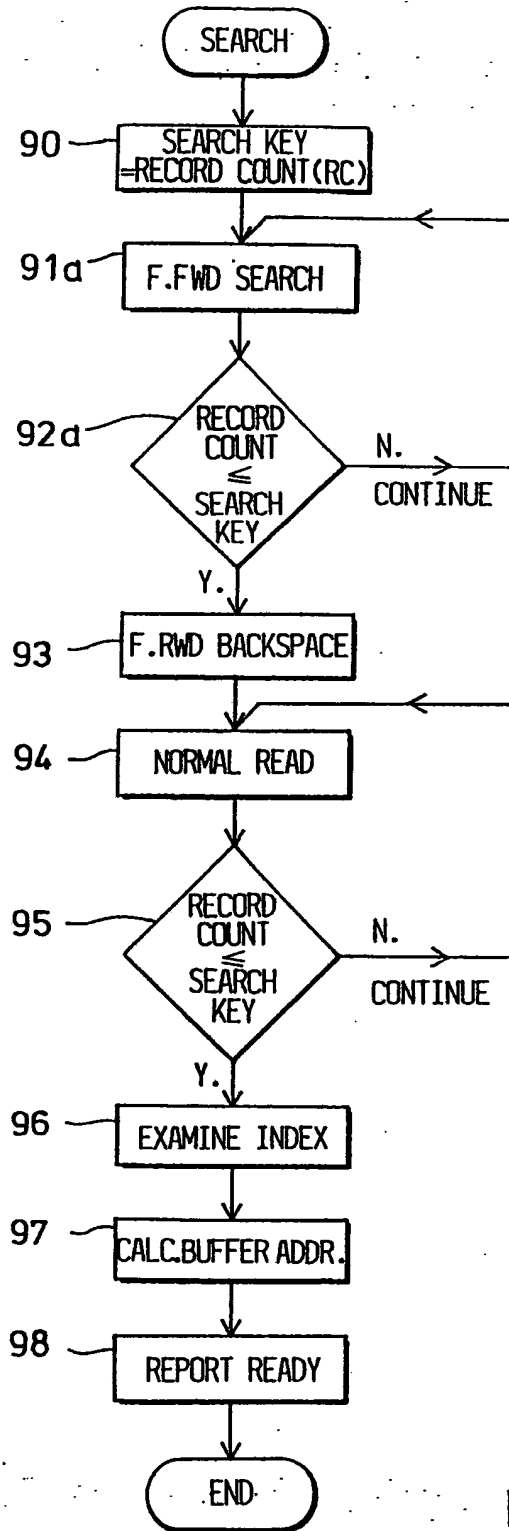


FIG 20 A

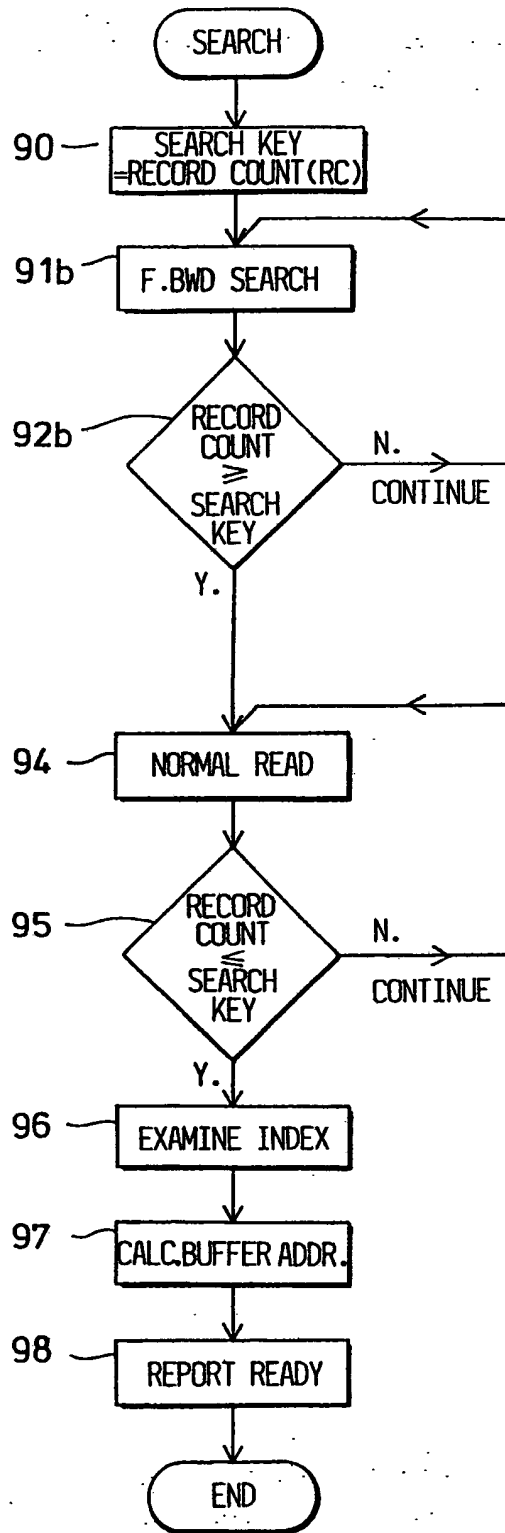



FIG 20B

INTERNATIONAL SEARCH REPORT

International Application No **PCT/GB 91/00082**

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ⁶		
According to International Patent Classification (IPC) or to both National Classification and IPC		
IPC ⁵ : G 11 B 20/00, G 11 B 20/12		
II. FIELDS SEARCHED		
Minimum Documentation Searched ⁷		
Classification System	Classification Symbols	
IPC ⁵	H 03 M, G 11 B	
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸		
III. DOCUMENTS CONSIDERED TO BE RELEVANT ⁹		
Category ¹⁰	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³
X	Hitachi Review, volume 36, no. 4, August 1987, (Tokyo, JP), K. Fujita et al.: "H-6485 magnetic tape subsystem", pages 193-200 see figure 6; page 197 ---	1,2,4,11
X	Hewlett-Packard Journal, volume 40, no. 3, June 1989, (Palo Alto, CA, US), M.J. Bianchi et al.: "Data compression in a half-inch reel-to-reel tape drive", pages 26-31 see page 30; left-hand column, lines 49-55 ---	1,3,11
Y	EP, A, 0327201 (HEWLETT-PACKARD LIMITED) 9 August 1989 see figures 2-9; column 4, line 8 - column 7, line 17	1,2,4
A	---	5,8,9
<p>¹⁰ Special categories of cited documents:</p> <ul style="list-style-type: none"> "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art. "Z" document member of the same patent family 		
IV. CERTIFICATION		
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report	
18th April 1991	07.06.91	
International Searching Authority	Signature of Authorized Officer	
EUROPEAN PATENT OFFICE	 Danielle van der Haas	

III. DOCUMENTS CONSIDERED TO BE RELEVANT (CONTINUED FROM THE SECOND SHEET)		
Category *	Citation of Document, " with indication, where appropriate, of the relevant passages	Relevant to Claim No.
Y	EP, A, 0324542 (HEWLETT-PACKARD COMPANY) 19 July 1989 see figures 1,2A,2B,2C; column 5, line 30; column 4, line 33	1,2,4
A	---	5,7
A	EP, A, 0323911 (HEWLETT-PACKARD LIMITED) 12 July 1989 ---	
A	EP, A, 0327188 (HEWLETT-PACKARD LIMITED) 9 August 1989 -----	

ANNEX TO THE INTERNATIONAL SEARCH REPORT
ON INTERNATIONAL PATENT APPLICATION NO.

GB 9100082

SA 43981

This annex lists the patent family members relating to the patent documents cited in the above-mentioned international search report. The members are as contained in the European Patent Office EDP file on 31/05/91. The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A- 0327201	09-08-89	JP-A- 1204274	16-08-89
EP-A- 0324542	19-07-89	US-A- 4891784 JP-A- 1295325	02-01-90 29-11-89
EP-A- 0323911	12-07-89	JP-A- 1204275	16-08-89
EP-A- 0327188	09-08-89	JP-A- 1223672	06-09-89

EPO FORM 60079

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82